

Experiences with the AEROnet/PSCN ATM Prototype

Richard S. Kurak¹, Anthony J. Lisotta², James D. McCabe², Alfred E. Nothaft², Kelly R. Russell²

Abstract

This paper discusses the experience gained by the AEROnet/PSCN networking team in deploying a prototype Asynchronous Transfer Mode (ATM) based network as part of the wide-area network for the Numerical Aerodynamic Simulation (NAS) Program at NASA Ames Research Center. The objectives of this prototype were to test concepts in using ATM over wide-area Internet Protocol (IP) networks and measure end-to-end system performance. This testbed showed that end-to-end ATM over a DS3 reaches approximately 80% of the throughput achieved from a FDDI to DS3 network. The 20% reduction in throughput can be attributed to the overhead associated with running ATM. As a result, we conclude that if the loss in capacity due to ATM overhead is balanced by the reduction in cost of ATM services, as compared to dedicated circuits, then ATM can be a viable alternative.

1. NASA Lewis Research Center

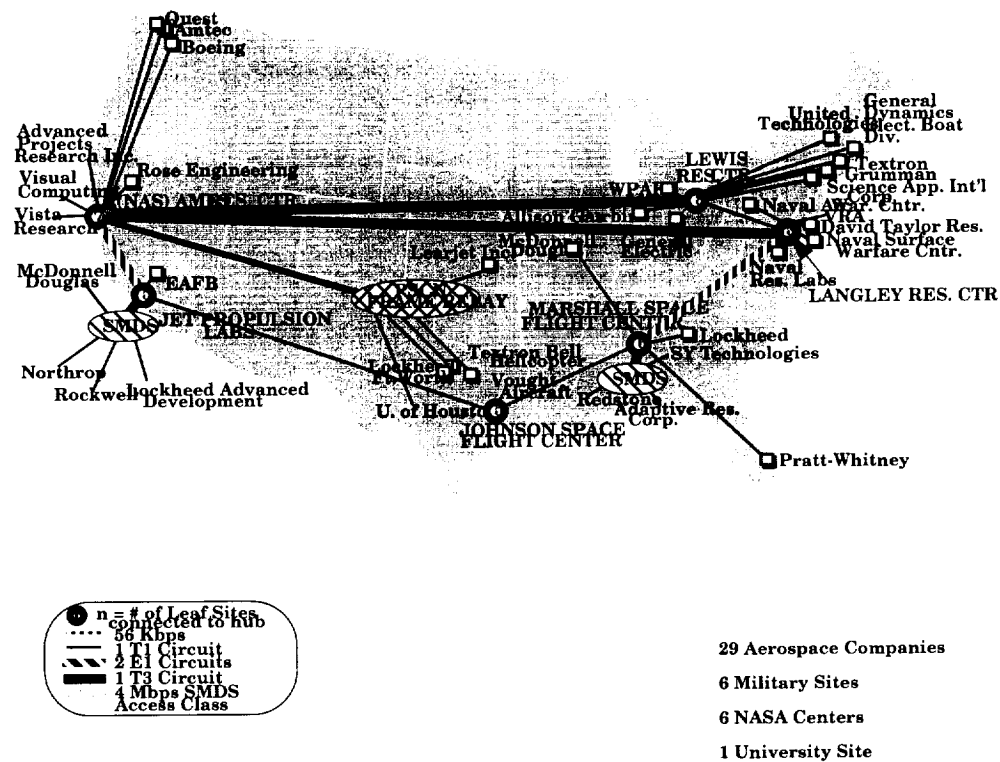
2. Computer Sciences Corporation, NASA Ames Research Center

1.0 Introduction

A prototype Asynchronous Transfer Mode (ATM) based network was deployed as part of the wide-area network for the Numerical Aerodynamic Simulation (NAS) Program at NASA Ames Research Center. This network was part of the NASA national aerospace network, AEROnet (see Figure 1). For more information about the AEROnet network, see the reference by Lisotta and McCabe^[1].

Objectives of the ATM prototype were to test concepts in using ATM over wide-area Internet Protocol (IP) networks, measure end-to-end system performance, and evaluate potential applications for wide-area ATM networks.

FIGURE 1. The NAS wide area network: AEROnet



The prototype was deployed from May - June, 1994, between the NAS Facility (Mountain View, CA), Lewis Research Center (Cleveland, OH), and Langley Research Center (Hampton, VA). Participants included the Program Support Communications (PSC) group at Marshall Space Flight

Center (Huntsville, AL), as well as engineers from the three sites on the prototype network.


It is expected that the results and experience gained from this prototype will be used in the deployment of a production quality ATM service for the AEROnet, and possibly other NASA networks. An important part of this ATM service is the configuration of the service at the demarcation points (where the wide-area service interfaces with local services). Our proposed configuration, termed here the *expected standard configuration*, is composed of three parts (see Figure 2).

The expected standard configuration allows the connection of IP routers, ATM switches, and end systems (workstations, supercomputers, mainframes, etc.) or any combination of these to the service demarcation ATM switch. Thus, the ATM service can be extended into the campus environment, or terminated at end systems. IP networks can also be joined to the ATM service.

2.0 Prototype Description

The ATM prototype was built upon the existing AEROnet infrastructure using dedicated DS3 circuits from the NAS facility to Langley Research Center (LaRC) and Lewis Research Center (LeRC) (see Figure 3). The purpose of this prototype was to gain experience in configuring ATM systems to work in an IP routed environment and to conduct experiments in the performance of ATM and IP cross-country networks. These experiments were conducted in four phases.

FIGURE 2. Expected Standard Configuration



Phase I connected the dedicated DS3 circuits mentioned above via IP routers to isolated FDDI networks at each of the centers. The purpose of this phase was to establish the baseline performance of an IP routed DS3

network. Phase II replaced the IP routers with ATM switches, with workstations connected directly to the switches. This basic configuration of ATM switches at NAS, LeRC, and LaRC, connected with dedicated DS3 circuits, formed the core upon which all remaining phases were based. In this phase the performance of an ATM DS3 network was determined. Phase III expanded the core ATM switched network to include the addition of ATM switches at LaRC and LeRC. This phase brought ATM into the local campus testbed networks. Phase IV, the final phase of the testbed, incorporated all of the components of the *expected standard configuration*. The purpose of this phase was to build the expected configurations for connection to wide-area ATM services.

2.1 Tests

Testing of the ATM prototype was organized into three groups: performance, functional, and application. Objectives of the testing were:

Performance Tests

- Determine maximum end-to-end performance of workstations directly connected to ATM switches over cross-country DS3 circuits.
- Compare performance of a cross-country DS3/ATM network connecting ATM switches at the three centers to a cross-country DS3/IP routed network connecting the same three centers.

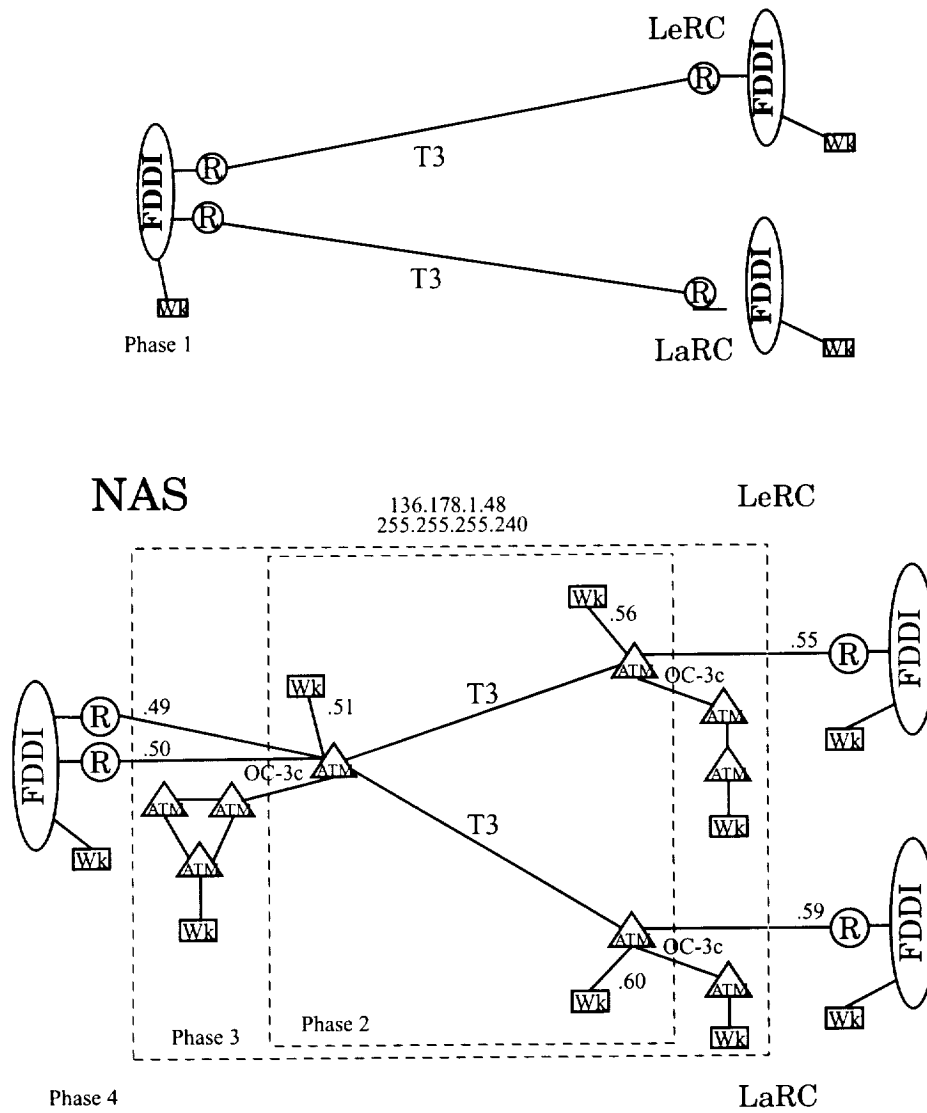
Functional Tests

- Utilize Permanent Virtual Connections(PVCs) and Switched Virtual Connections (SVCs) across the ATM network.
- Increase the number of switches and connections by connecting LeRC and LaRC ATM testbeds to the AEROnet prototype.
- Connect ATM switches to FDDI testbed networks at NAS, LaRC, and LeRC through the use of IP routers.
- Evaluate OSPF and RIP in all phases.

Applications Tests

- Test videoconferencing using ATM.
- Run standard applications (telnet, FTP, rlogin, etc.) over the ATM network.

FIGURE 3. ATM Prototype Network



3.0 Results

Performance tests consisted of a series of transfers using *ttcp* to measure data transfer rates. The performance measuring tests varied the number of tcp streams as well as the data buffering size. Socket buffer size was set to the maximum size allowed by the operating system configuration, 64 KB for SGIs and 48 KB for Sun workstations.

TABLE 1. Data Summary:

	DS3/FDDI	ATM WKS	DS3/ATM
Maximum Data Rate	5.2406 MB/s	3.47726 MB/s	4.22527 MB/s
Maximum Bit Rate	41.924 Mb/s	27.818 Mb/s	33.802 Mb/s
Overhead per byte ^a	0.0117 bytes	0.117 bytes	0.117 bytes

a. Assume 4096 bytes of data per IP packet.

3.1 Baseline Results

The initial phase of the tests was to measure the baseline performance of a routed DS3/FDDI network. The main goal of these tests were to determine if the routers could forward traffic at sustained DS3 rates (~44 Mb/s). The graphs shown in Figures 4 through 7 are the results of the tests between the NAS facility & LARC and between the NAS facility & LERC. The workstation at the NAS facility was an SGI 4D/320, the workstations at LARC and LERC were both Sun Sparc 10s. Each workstations was connected to a dedicated FDDI ring as shown in Figure 3.

Multiple streams were used to overcome limitations of the workstations and operating systems to fill the entire DS3. All tests used TCP transmissions because this best simulates the types of data transfers commonly used by scientists on the AEROnet network.

In the graphs shown in Figures 4 through 7 above, note that transfers with 1, 2 and 4 TCP streams are close to being multiples of each other (i.e. Data rate using 2 streams is ~2 * data rate using 1 stream, and data rate using 4 streams is ~4 * data rate using 2 streams). This suggests that these data rates are limited by the workstations' operating systems, not the network. However, once we approach 8 & 9 TCP streams the data rate begins to level out and little to no increases are experienced. This suggests that the network was the limiting factor and thus this is approximately the maximum sustained load that can be put on the network.

The data rates recorded do not include the IP or TCP overhead, this overhead would account for an additional 40 bytes/pkt.

FIGURE 4. NAS to LaRC Straight DS3




FIGURE 5. LaRC to NAS Straight DS3




FIGURE 6. NAS to LeRC Straight DS3





FIGURE 7. LeRC to NAS Straight DS3



3.2 Phase 2

The graphs show in Figures 8 through 10 are the results from performance tests between workstations connected directly to the ATM network by 100 Mb/s TAXI interfaces as shown in the phase two portion of Figure 3.

In this series of tests, tests using 1 and 2 TCP streams appear to be bound by the workstations not the network. However tests using 4 or more streams appear to be leveling off, we believe this is due more to the network interface card than the network itself.

FIGURE 8. NAS to LaRC workstation to workstation with ATM/DS3




FIGURE 9. NAS to LeRC workstation to workstation with ATM/DS3





FIGURE 10. LeRC to NAS workstation to workstation with ATM/DS3



3.3 Phase 3

Due to differing revisions on the local area ATM testbeds, no performance numbers were attained.

3.4 Phase 4

The final phase of the tests was aimed at measuring the maximum throughput capabilities of an ATM network utilizing DS3 links. Again, TCP streams were used because they most closely represent the type of traffic that these networks would carry in the AEROnet environment. Unfortunately the link between the NAS facility and Langley Research Center could not be stabilized for these tests and thus data is not available for that portion of the network, see Section 3.1 for further details.

The graphs show in Figures 11 and 12 are results from tests using multiple TCP streams to transfer data across the ATM/Router network. The results from these tests closely resemble the results from the tests utilizing a straight DS3/FDDI for 1, 2 and 4 streams. However, for 8 or more streams, the maximum sustained data rates using DS3/ATM are 0.05-1MB/s less than the equivalent number of streams over the DS3/FDDI network.

FIGURE 11. NAS to LeRC workstation to workstation through IP routers and with ATM/DS3





FIGURE 12. LeRC to NAS workstation to workstation through IP routers and with ATM/DS3



3.5 Experiences

Not unexpectedly, there were times during the deployment and testing of the ATM prototype network when the prototype exceeded expectations, as well as times when the network did not work or performed poorly. This section describes some of these experiences.

Permanent and Switched Virtual Connections (PVCs and SVCs)

The ATM switches were initially configured as a fully meshed PVC network. This configuration process, which occurred at multiple centers at the same time, became unwieldy and difficult to maintain. It was decided early in the configuration process that PVC provisioning was not feasible in a dynamic environment. As a result, SVCs were used instead of PVCs, where possible.

Data Service Units (DSUs)

ATM-capable DSUs were used to connect the IP routers to ATM switches, using HSSI interfaces at the routers and DS3 interfaces at the switches. Instabilities with the DSUs, combined with routing protocol problems, resulted in a partial deployment of Phase IV. Phase IV performance data

was collected only on the NAS-LeRC portion of the prototype. The instabilities were due to an interoperability problem between the DSUs and IP routers.

As both the router and DSU vendors did not have a stable ATM Data Exchange Interface (DXI) at the time of prototype deployment, we used the Frame Relay DXI instead.

Video over ATM

As part of the applications testing during Phase III, ATM video adapters from Fore Systems (AVA-200) were deployed at each of the three centers, and were directly connected to the ATM switches at each site. These adapters provided an ATM interface for various audio and video equipment. For the Phase III video trials, we used NTSC video without compression. Each site produced a video stream to the ATM network, where PVCs were configured to the originating site as well as each of the other centers. This resulted in a three-way simultaneous video session, where each site could view and interact with all sites. At the time of the deployment of the video adapters, the AVA-200s were in the alpha stage of testing.

Quality of the video over the prototype network was exceptional. There were three simultaneous video sessions, each running at 12 frames/second and window sizes of approximately 320 x 240 pixels. While running at the maximum data rates, there was little to no distortion in the images, even during periods of rapid movement. It was understood at the beginning of this phase of testing that the video adapters were alpha versions, and that configuring multiple audio streams would not be possible. For our tests, we were able to achieve audio in one direction only.

Protocols

During this prototype, TCP/IP was used for the transport and network layer protocols. For the routing across the network, two protocols were tested, RIP and OSPF. The workstations and switches used SVCs to communicate with each other, and PVCs were setup for routers to communicate with each other. Although the routers were able to pass IP traffic across the ATM network, we were unsuccessful in getting routers to communicate with either switches or hosts. We believe this is due to a problem with the address resolution schemes.

As for routing protocols, RIP functioned as expected and no problems were encountered with this protocol. OSPF on the other hand, did not function. There was a problem with the routers establishing full adjacencies with each other. During the test the cause of the problem was not known; however, after extensive research after the trial, it was deter-

mined that there was a problem with the size of the OSPF MTU. This caused some of the OSPF packets to be corrupted. As a result, OSPF was not used in the testing.

4.0 Conclusions

The initial goal of the ATM prototype was to compare an ATM/DS3 network to a standard IP routed network using point-to-point DS3 circuits, and to determine the maximum end-to-end performance of an ATM network. The DS3/ATM and DS3/FDDI performance results reached the approximate theoretical maximum for each technology. If the loss in capacity due to ATM overhead is balanced by the reduction in cost of ATM services, as compared to dedicated circuits, then we believe that ATM can be a viable alternative. In the process of evaluating the ATM prototype, we have also discovered the following:

- PVC provisioning and maintenance becomes unwieldy and SVCs need to be supported by all end equipment.
- The use of DSUs as an interface into an ATM service adds complexity, and requires the use of PVCs.
- Video over native ATM was very good.
- Vendors must increase performance of ATM adapter cards.
- Multicast functionality will become more important as tele* (teleseminar, telelearning, teleconference, etc.) applications become prevalent.

5.0 Future Directions

Future tests of ATM over the wide area can be approached from several directions. As work already has been done, ATM experiments can focus into any of several areas of the network. Future testing can highlight any one or more of hardware, software, applications, or the entire ATM environment.

Hardware

The ATM video adapter from Fore Systems (AVA-200) shows promise as a hardware platform for desktop videoconferencing. However, at the time of initial testing, this product was in alpha state. It needs to undergo further testing to determine what its practical limitations are. Utilizing the DS3 circuit between sites, the practical number of concurrent audio and video streams the device can support over the available bandwidth needs to be determined. The AVA environment can be further studied for bandwidth impacts of one-to-many (presentation), many-to-one (closed circuit video), one-to-many video with many-to-many audio (distance

learning) and many-to-many (full video conferencing) sessions. Each of these can be further taken to account for multiple session groupings on the same network. If full distribution is not required, how many simultaneous but diverse sessions can then be taking place across one or more ATM switches.

The tests performed for this paper used only Fore systems switches. By having a homogenous environment, issues about interoperability are not involved. Future tests will expand the testbed to include other vendors' ATM switches. The test suite used in this paper will be duplicated using other vendors' equipment. Additionally, multiple vendors' equipment will be integrated into a single network to determine performance in conjunction with interoperability.

Software (Circuit Maintenance)

Different parts of the experiments already conducted have had requirements for either using PVCs or SVCs. While some equipment requires the use of a PVC, the administrative overhead associated with maintaining such a network rises exponentially with the number of stations. SVCs do not have administrative overhead, but incur small amounts of overhead in the call setup and tear down. Further testing can be used to determine under what conditions PVCs or SVCs should be used. The testing on each will take into account additional latency added by call setup and teardown for an overall performance figure.

Applications

While the ATM environment shows great promise from a strict networking perspective, it will be user applications that will make ATM a viable alternative to current network methodologies. One such application that could take advantage of the bandwidth and latency characteristics of ATM is the distributed Virtual Wind Tunnel (dVWT). In the dVWT, 3-dimensional models are simulated in a computational wind tunnel. The researchers can then enter the windtunnel during the test to look at or adjust the model from different angles and directions. The model is displayed via a 3-dimensional viewer that allows the researcher to enter the simulation. The dVWT is at the high end of applications that can benefit from ATM. Any application that requires high speed and high bandwidth with strict latency requirements will benefit. Many applications in the Computational Fluid Dynamics (CFD) arena are potential candidates. Such applications can be characterized by high computational requirements with large amounts of data. The results, although in the past have been data sets of numbers (which although large did not have strict latency requirements), if displayed in real-time using video would make current FDDI and Ethernet networks almost impossible to use. By using ATM, the tests being run can be adjusted interactively to produce better

results in a fraction of the time. Further testing can show the viability and potential gains of utilizing ATM to create such and applications environment, as well as possibly quantizing some of the benefits gained.

ATM for Workstation Clusters


Experiments are currently being done at LeRC to determine the applicability of clustering workstations across various media including ATM. The tests include tightly coupled and loosely coupled problems. Results center on various performance characteristics. These tests can be extended across the wide area in order to determine what, if any, problems are incurred by the cross-country propagation delays of a wide-area network. Such an environment would allow under-used resources at one location (possible due to off-peak hours in a different time-zone) to be applied to applications at another site.

Media Performance

One of the benefits of ATM is the ability to get and maintain strict latency and bandwidth requirements. If a store and forward device (such as a router) is placed into the stream, the advantages of end-to-end ATM are lost. More problems could be incurred when on the other side of the router a different media is added. Both ATM and FDDI offer bandwidths of 100 Mbps. Tests that make direct comparisons of 100 Mbps pipes could show the benefits and problems of using either type of media. Each type will have its place in networking for some time to come, if each is utilized under the proper circumstances. Future testing should show the basic criterion for each of these media. When ATM and FDDI are combined in the same network, a different set of issues could arise. Performance figures could be arrived at to show some of the advantages or pitfalls of such a network as well as whether the choke point is the media or the interconnection device.

6.0 References

- [1] Lisotta, Anthony and James D. McCabe. "AEROnet: The NASA Aerospace Network", NASA Ames Research Center, 1994.
- [2] ATM User-Network Interface Specification, Version 3.0. The ATM Forum, PTR Prentice Hall, New Jersey, 1993.

	<h1>NAS TECHNICAL REPORT</h1>	
	<p>Title: Experiences with the AEROnet/PSCN ATM Prototype</p>	
	<p>Author(s): Richard S. Kurak^a, Anthony J. Lisotta^b, James D. McCabe^b, Alfred E. Nothaft^b, Kelly R. Russell^b</p>	
	<p>Reviewers: "I have carefully and thoroughly reviewed this technical report. I have worked with the author(s) to ensure clarity of presentation and technical accuracy. I take personal responsibility for the quality of this document."</p>	
<p>Two reviewers must sign.</p>	<p>Signed: _____</p> <p>Name: _____</p> <p>Signed: _____</p> <p>Name: _____</p>	
<p>After approval, assign NAS Report number.</p>	<p>Branch Chief:</p> <p>Approved: _____</p>	
<p>Date:</p>	<p>NAS ReportNumber:</p>	

a. NASA Lewis Research Center

b. Computer Sciences Corporation, NASA Ames Research Center